UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

| APPLICATION NO. | FILING DATE | FIRST NAMED INVENTOR | ATTORNEY DOCKET NO. | CONFIRMATION NO. |
|---|---|---|---|---|
| 10/646,851 | 08/22/2003 | Rui Liang | 10830.0097NP | 1500 |

27927 7590 10/25/2010

RICHARD AUCHTERLONIE
NOVAK DRUCE & QUIGG, LLP
1000 LOUISIANA
53RD FLOOR
HOUSTON, TX 77002

| EXAMINER | |
|---|---|
| FLYNN, NATHAN J | |

| ART UNIT | PAPER NUMBER |
|---|---|
| 2468 | |

| MAIL DATE | DELIVERY MODE |
|---|---|
| 10/25/2010 | PAPER |

**Please find below and/or attached an Office communication concerning this application or proceeding.**

The time period for reply, if any, is set in the attached communication.

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE BOARD OF PATENT APPEALS AND INTERFERENCES

*Ex parte* RUI LIANG, JIANNAN ZHENG, MARK K. KU, XIAOYE JIANG, and SORIN FAIBISH

Appeal 2009-006658
Application 10/646,851
Technology Center 2400

Before JOHN A. JEFFERY, JEAN R. HOMERE, and JAMES R. HUGHES, *Administrative Patent Judges*.

JEFFERY, *Administrative Patent Judge*.

DECISION ON APPEAL[1]

Appellants appeal under 35 U.S.C. § 134(a) from the Examiner's rejection of claims 7-11, 19-23, 27-36, and 40-49. We have jurisdiction under 35 U.S.C. § 6(b). We affirm-in-part.

[1] The two-month time period for filing an appeal or commencing a civil action, as recited in 37 C.F.R. § 1.304, or for filing a request for rehearing, as recited in 37 C.F.R. § 41.52, begins to run from the "MAIL DATE" (paper delivery mode) or the "NOTIFICATION DATE" (electronic delivery mode) shown on the PTOL-90A cover letter attached to this decision.

STATEMENT OF THE CASE

Appellants invented a storage object (e.g., a virtual disk drive or logical volume) contained in a UNIX compatible file so that it can be exported using various protocols and shared among UNIX and Windows clients or servers. *See generally* Abstract; Spec. 3-5; Figs. 1, 2, 5. Claim 27 is illustrative with key disputed limitations emphasized:

27. A network file server comprising:

data storage;

an interface for coupling the data storage to a data network; and

at least one processor programmed for permitting clients in the data network to access the data storage in accordance with a plurality of access protocols;

the data storage containing at least one file for storing file attributes and for storing metadata defining a virtual direct access storage device and for storing data of the virtual direct access storage device;

the access protocols including at least one block level access protocol for access to the virtual direct access storage device by accessing the metadata and data of the virtual direct access storage device; and

the access protocols including at least one file access protocol for accessing said at least one file;

wherein the metadata includes attributes of the virtual direct access storage device, and the attributes of the virtual direct access storage device and the data of the virtual direct access storage device are stored together in a single file in a file system; and

wherein the attributes of the virtual direct access storage device include a *specification of an internal organization of the virtual direct access storage device for mapping of the data of the virtual direct access*

*storage device from the single file to the data storage, and the specification of the internal organization of the virtual direct access storage device is stored in the single file.*

The Examiner relies on the following as evidence of unpatentability:

| | | |
|---|---|---|
| Hitz | US 5,819,292 | Oct. 6, 1998 |
| Bolosky | US 2001/0052021 A1 | Dec. 13, 2001 |
| Lefebvre | US 2002/0010665 A1 | Jan. 24, 2002 |
| Busser | US 2002/0095616 A1 | July 18, 2002 |
| Baweja | US 6,564,229 B1 | May 13, 2003 |
| Hashemi | US 6,934,804 B2 | Aug. 23, 2005 (filed May 28, 2002) |
| Chen ("Chen '553") | US 7,010,553 B2 | Mar. 7, 2006 (filed Mar. 19, 2002) |
| Chen ("Chen '509") | US 7,076,509 B1 | July 11, 2006 (filed Mar. 21, 2003) |
| Rajan | US 7,107,385 B2 | Sept. 12, 2006 (filed Aug. 9, 2002) |

THE REJECTIONS[2]

1. The Examiner rejected claims 27, 30, and 34-36 under 35 U.S.C. § 102(e) as anticipated by Chen '509. Ans. 4-7.[3]
2. The Examiner rejected claims 7-11 and 19-23 under 35 U.S.C. § 103(a) as unpatentable over Chen '509, Baweja, and Lefebvre. Ans. 7-12.

[2] Since the Examiner withdrew a rejection under § 112 (Ans. 2, 19), that rejection is not before us.
[3] Throughout this opinion, we refer to (1) the Appeal Brief filed May 23, 2008; (2) the Examiner's Answer mailed August 4, 2008; and (3) the Reply Brief filed October 5, 2008.

3. The Examiner rejected claims 28 and 29 under 35 U.S.C. § 103(a) as unpatentable over Chen '509, Busser, and Hashemi. Ans. 12-13.
4. The Examiner rejected claims 33, 40, and 45-48 under 35 U.S.C. § 103(a) as unpatentable over Chen '509 and Chen '553. Ans. 13-16.
5. The Examiner rejected claims 41 and 42 under 35 U.S.C. § 103(a) as unpatentable over Chen '509, Chen '553, Busser, and Hashemi. Ans. 17-18.
6. The Examiner rejected claims 31, 32, 43, 44, and 49 under 35 U.S.C. § 103(a) as unpatentable over Chen '509, Chen '553, and Bolosky. Ans. 18.

## THE ANTICIPATION REJECTION

Regarding representative claim 27, the Examiner finds that Chen '509 discloses a network file server with every recited feature including data storage containing a file for storing (1) file attributes, and (2) metadata including attributes of the virtual direct access storage device that is stored together in a single file with associated data as claimed. Ans. 4-7. The Examiner cites the incorporated Rajan reference to support the contention that the attributes include a specification of an internal organization of the virtual direct access storage device for mapping the device's data from the single device to the data storage as claimed. Ans. 19-20.

Appellants disagree with this contention, since the RAID redundancy or striping in Chen '509 and Rajan is said to be a property of the volume upon which the file system is built, rather than an attribute of the disclosed vdisk file. App. Br. 22-23. Although Appellants acknowledge Rajan's attributes inode and stream_dir inode metadata section in connection with

vdisk representations, Appellants nonetheless contend that there is no internal organization specified for mapping the virtual direct access storage device's data from the single file to the data storage as claimed. App. Br. 23-24; Reply Br. 3-4. The issue before us, then, is as follows:

## ISSUE

Under § 102, has the Examiner erred in rejecting claim 27 by finding that Chen '509 discloses data storage containing a file for storing (1) file attributes, and (2) metadata including attributes of the virtual direct access storage device that is stored together in a single file with associated data, where the attributes include a specification of an internal organization of the virtual direct access storage device for mapping the device's data from the single device to the data storage?

## FINDINGS OF FACT (FF)

1. Appellants indicate that the claimed attributes of the virtual direct access storage device correspond to the storage object attributes 86 in Figure 5 of the present application. App. Br. 10-11 ("Summary of Claimed Subject Matter" section).

2. The storage object attributes 86 of the present application include a (1) storage object type (e.g., disk, logical volume, etc.); (2) storage capacity; (3) security information; (4) storage object location; and (5) an internal organization of the storage object (e.g., level of redundancy in a disk drive array (RAID level) and a striping scheme). Spec. 11:1-18; Fig. 5.

3. Figure 5 of the present application includes the textual label "INTERNAL ORGANIZATION (E.G., RAID LEVEL, STRIPING)" underneath the label "STORAGE OBJECT ATTRIBUTES."

4. "The specified internal organization of the storage object could be used as a guide or specification for mapping of the data storage area 87 of the container file 87 [sic] to storage in the cached disk array . . . ." Spec. 11:16-18.

5. Chen '509 incorporates Application Serial Number 10/216,453 (now U.S. Pat. 7,107,385) (Rajan) by reference. Chen '509, col. 1, ll. 6-11.

6. Chen '509's multi-protocol storage appliance 500 includes a storage operating system 600 that implements a file system, and provides for storage virtualization and support for virtual disks. The storage appliance enables file and block protocol access to stored information, and exports disks to clients by creating logical unit numbers ("luns") or vdisk objects ("vdisks"). Chen '509, col. 7, ll. 58-62; col. 8, ll. 4-31; Fig. 5.

7. A vdisk is a special file type implemented by the virtualization system and translated into an emulated disk as viewed by clients. A vdisk is a multi-inode object comprising (1) a special file inode, and (2) a set of stream inodes that are managed as a single encapsulated storage object within a file system. The vdisk encapsulates security, management, and addressing (export) information, and comprises (1) a data section, and (2) metadata section(s) that may be stored in streams associated with the data section. Chen '509, col. 8, ll. 31-33; col. 10, ll. 43-60.

8. Information storage on appliance 500 is preferably implemented as one or more storage volumes (e.g., VOL 1-2 550) comprising a cluster of physical storage disks 530 defining an overall logical arrangement of disk

space. The disks within a volume are typically organized according to a RAID implementation (e.g., RAID groups 540, 542, 544). Chen '509, col. 10, ll. 17-42; Fig. 5.

9. Rajan discloses a multi-protocol storage appliance 100 similar to that shown in Chen '509. *Compare* Rajan, Fig. 1 *with* Chen '509, Fig. 1.

10. Rajan's virtualization system 300 includes a vdisk module 330 layered on file system 320 to enable access via administrative interfaces. The file system's on-disk representation uses inodes to describe the files, as well as metadata files that include an inode file. Rajan, col. 8, l. 12 – col. 9, l. 12; Fig. 3.

11. Rajan's on-disk inode includes (1) a metadata section 410, and (2) a data section 450. The metadata section (1) describes the file by including information pertaining to the file type, size, ownership, etc., and (2) includes an xinode field 430 containing a pointer referencing another on-disk inode structure. Rajan, col. 9, ll. 20-39; Fig. 4.

12. In one implementation, Rajan's visualization technique provides an on-disk representation of vdisk 322 sorted on the multi-protocol storage appliance. This representation includes vdisk (lun) and stream (attributes) inodes. Stream_dir inode 520 comprises a metadata section 522 including (1) a type (stream dir) field 524, and (2) an xinode field 525 that references another on-disk inode structure. The stream_dir inode also includes a data section containing a pointer 528 referencing a stream directory data block 530 associated with the vdisk. The stream directory block contains mapping information such as an inode number and a pointer referencing attributes (stream) inode 540. Rajan, col. 12, l. 48 – col. 13, l. 30; Fig. 5.

13. Rajan's system manages a vdisk (data and attributes) as a single encapsulated unit within the file system. The binding between a vdisk's data (file inode) and attributes (stream inode) creates a single encapsulated object that is "self-describing" in that it contains all information necessary to, e.g., access that object. Rajan, col. 14, ll. 35-62.

## ANALYSIS

Based on the record before us, we find no error in the Examiner's anticipation rejection of representative claim 27. Claim 27 calls for, in pertinent part, data storage containing a file for storing (1) file attributes, and (2) metadata including attributes of the virtual direct access storage device that is stored together in a single file with associated data, where the attributes include a specification of an internal organization of the virtual direct access storage device for mapping the device's data from the single device to the data storage.

At the outset, we note that the particulars of the attributes noted above in the last clause of claim 27 merely describe the content of the information and do not further limit the claimed invention functionally. Simply put, the recited "specification of an internal organization of the virtual direct access storage device" is merely informational content and therefore non-functional descriptive material. Such non-functional descriptive material does not patentably distinguish over prior art that otherwise renders the claim unpatentable.[4]

[4] *See In re Ngai*, 367 F.3d 1336, 1339 (Fed. Cir. 2004); *see also Ex parte Nehls*, 88 USPQ2d 1883, 1887-89 (BPAI 2008) (precedential) (discussing cases pertaining to non-functional descriptive material).

Nevertheless, even if the recited specification were functional (which it is not), we are still not persuaded of error in the Examiner's position. First, as the Examiner indicates (Ans. 19-20), the recited "specification of an internal organization of the virtual direct access storage device" is hardly limited to RAID levels or striping, for as Appellants' own disclosure indicates, these schemes are merely exemplary "internal organizations." FF 2-3. Indeed, claim differentiation principles alone evidence this fact: otherwise, dependent claims claim 28 and 29 (limiting the specified internal organization to including a RAID level and striping pattern, respectively) would be superfluous.[5] In any event, that Appellants' Figure 5 actually uses the term "e.g." in connection these schemes as exemplary internal organization attributes of a storage object (FF 3)—attributes that Appellants equate to the recited attributes (FF 1)—only bolsters this conclusion.

Based on the scope and breadth of the term "internal organization," we therefore are not persuaded of error in the Examiner's reliance on the extensive mapping and reference information contained within the various inodes in connection with Rajan's[6] vdisk functionality in Figure 5. *See* FF 10-13. We see no reason why this information does not constitute, at least in part, a specification of an internal organization of the virtual direct access

[5] "The doctrine of claim differentiation creates a presumption that each claim in a patent has a different scope . . . . The difference in meaning and scope between claims is presumed to be significant to the extent that the absence of such difference in meaning and scope would make a claim superfluous." *Free Motion Fitness, Inc. v. Cybex Int'l, Inc.*, 423 F.3d 1343, 1351 (Fed. Cir. 2005) (internal quotation marks and citations omitted).
[6] The relied-upon anticipatory reference to Chen '509 incorporates Rajan by reference. FF 5.

storage device, given the scope and breadth of the limitation. Appellants' arguments to the contrary (App. Br. 22-24; Reply Br. 3-4) are simply not commensurate with the scope of the claim.

And as the Examiner indicates (Ans. 20), that claim 27 recites that this internal organization specification is "*for* mapping of the data of the virtual direct access storage device from the single file to the data storage" (emphasis added) merely recites an intended use of the specification. Where, as here, the prior art specification is capable of performing this intended use, it meets the disputed limitation. *See In re Schreiber*, 128 F.3d 1473, 1477-78 (Fed. Cir. 1997).

Despite Appellants' arguments to the contrary (Reply Br. 4-8), Appellants have simply not shown error in the Examiner's position (Ans. 20) that the information provided in Rajan's Figure 5 could be used, at least in part, to map data from the single file to the data storage device as claimed. We reach this conclusion noting the extensive mapping and reference information contained within the various inodes in connection with Rajan's vdisk functionality. *See* FF 10-13.

Nor are we persuaded of error in the Examiner's reliance on Chen '509 (Ans. 4-5) for it incorporates Rajan's disclosure by reference (FF 5), and both references disclose commensurate multi-protocol storage appliance vdisk functionality. *See* FF 6-13. Although the physical disks in connection with each reference's storage appliance are implemented as storage volumes as Appellants contend (App. Br. 23; FF 8-9), Appellants' arguments are

simply not commensurate with the scope of the claim which, as noted previously, does not preclude the Examiner's reliance on the vdisk "internal organization" specified via Rajan's visualization functionality. *See* FF 10-13.

Appellants also contend that Chen '509 and Rajan teach away from including in a container file together with data of the virtual direct access storage device, RAID redundancy, striping, and related information. App. Br. 24. But this argument, too, is not commensurate with the scope of the claim and, in any event, teaching away arguments are irrelevant to anticipation. *Leggett & Platt, Inc. v. VUTEk, Inc.*, 537 F.3d 1349, 1356 (Fed. Cir. 2008) (citation omitted). Nevertheless, we find no error in the Examiner's position that Rajan stores the recited information in a single file indicating an internal organization of the virtual direct access storage device. Ans. 19.

We are therefore not persuaded that the Examiner erred in rejecting representative claim 27, and claims 30 and 34-36 not separately argued.

## THE OBVIOUSNESS REJECTION OVER CHEN '509, BAWEJA, AND LEFEBVRE

Regarding representative claim 8, the Examiner finds that Chen '509 discloses every recited feature except for the client pausing writing data to the storage object after a commit operation, and during the pause, the client initiating file copying. The Examiner, however, cites Baweja to cure this deficiency. Ans. 7-8, 10-11. The Examiner also acknowledges that Chen '509 and Baweja collectively fail to teach that the client initiates file copying

by sending a command over a second “concurrent” TCP/IP connection, but cites Lefebvre to cure that deficiency in concluding the claim would have been obvious. Ans. 11.

Appellants argue that the cited prior art does not teach or suggest a client using block level access protocol to write data to a storage object in the first file server over a first TCP/IP connection concurrent with the first file server replicating a snapshot copy of the file to the second file server, such that the client initiates this replication by sending a command to the first file server over a second TCP/IP connection while the client pauses writing data to the first file server as claimed. App. Br. 25-33; Reply Br. 8-10. The issue before us, then, is as follows:

## ISSUE

Under § 103, has the Examiner erred in rejecting claim 8 by finding that Chen ‘509, Baweja, and Lefebvre collectively would have taught or suggested the client using block level access protocol to write data to a storage object in the first file server over a first TCP/IP connection concurrent with the first file server replicating a snapshot copy of the file to the second file server, such that the client initiates this replication by sending a command to the first file server over a second TCP/IP connection while the client pauses writing data to the first file server?

## ADDITIONAL FINDINGS OF FACT

14. Chen ‘509 incorporates U.S. Patent 5,819,292 (Hitz) by reference. Chen ‘509, col. 3, ll. 44-52.

15. Chen '509 restores a vdisk from a snapshot without needing to copy every block or inode within the active file system. A vdisk restore process (1) duplicates the vdisk's inode, and (2) reconciles the blocks of the twin and snapshot inodes. Chen '509, Abstract; col. 15, l. 45 – col. 16, l. 33; Fig. 8.

16. Hitz notes that backups can be performed from a file system "clone" (i.e., a read-only copy of the file system on disk) instead of from the active file system. This allows the file server to remain online during backup. Hitz, col. 1, ll. 27-33.

17. Baweja's move/copy interface has a pause feature (e.g., via button 140 in Figure 1a) allowing the user to pause, and later resume, a file move or copy command via button 160 in Figure 1b. Baweja, Abstract; col. 3, ll. 4-57; Figs. 1a-1b.

18. Lefebvre's database management system 140 supports snapshot replication where data snapshots can be taken upon request. Lefebvre, ¶¶ 0039, 0085-86; Figs. 1, 5.

## ANALYSIS

We will not sustain the Examiner's obviousness rejection of independent claim 8. Claim 8 recites, in pertinent part, that the client uses block level access protocol to write data to a storage object in the first file server over a first TCP/IP connection *concurrent with* the first file server replicating a snapshot copy of the file to the second file server, such that *the client* initiates this replication by sending a command to the first file server over a second TCP/IP connection *while the client pauses writing data to the first file server.*

We emphasize these key disputed temporal limitations, for on the record before us, we cannot say that the cited prior art reasonably teaches or suggests these limitations when considering the claim as a whole. First, the Examiner cites Hitz (Ans. 8, 22) (which Chen '509 incorporates by reference (FF 14)) which teaches that backups can be performed on file system "clones" instead of active file systems, and therefore enable the server to remain online during backup. FF 16. Although somewhat general, this teaching nonetheless reasonably suggests that file server operations (e.g., reading, writing) on the "active file system" can occur concurrently with backing up or replicating its "cloned" copy.

In light of this teaching, and given Chen '509's snapshot restoration techniques (FF 15), we find reasonable the Examiner's position that these teachings collectively would have at least suggested the recited snapshot copy replication concurrent with the client writing data to the file server.

But we fail to see how the prior art reasonably teaches or suggests the other key temporal limitation of claim 8, namely that the client initiates the replication by sending a command to the first file server over a second TCP/IP connection *while the client pauses writing data to the first file server*.

Leaving aside the question of whether using a different TCP/IP connection to achieve this end would have been obvious in light of Lefebvre and the cited extrinsic evidence as the Examiner asserts (Ans. 21), we still find the Examiner's reliance on Baweja problematic regarding the specific temporal requirements of the client's actions noted above.

Baweja provides an interface allowing the user to pause, and later resume, file move and copy operations. FF 17. To be sure, this somewhat general teaching provides a rational basis for the Examiner's position regarding the obviousness of the client's pausing writing data to the first file server as claimed. But we still fail to see how Baweja teaches that the client initiates the first server's replicating the file's snapshot copy by sending a command over a different TCP/IP connection *during this pause*—a crucial temporal requirement of claim 8. And while we agree with the Examiner (Ans. 11) that Lefebvre suggests that clients can request snapshot replication from a server (FF 18), there is nothing on this record to suggest that such a client request—in the form of a command sent over a different TCP/IP connection or otherwise—would be made *while the client pauses writing data to the server* as claimed.

We are therefore persuaded that the Examiner erred in rejecting (1) independent claim 8; (2) independent claim 20 which recites commensurate limitations; and (3) dependent claims 7, 9-11, 19, and 21-23 for similar reasons. Since our decision is dispositive of our reversal of the Examiner's rejection of these claims, we need not address Appellants' other arguments pertaining to (1) the combinability of Chen '509, Baweja, and Lefebvre (App. Br. 32-33), and (2) dependent claims 7 and 19 (App. Br. 33-34; Reply Br. 9-10).

THE OBVIOUSNESS REJECTION OVER CHEN '509, BUSSER, AND HASHEMI

The Examiner finds that Chen '509 does not disclose that the specified internal organization includes a RAID level (claim 28) and striping

pattern (claim 29), but cites Busser and Hashemi to cure these respective deficiencies in concluding the claims would have been obvious. Ans. 12-13.

Appellants argue that Chen '509 and Rajan teach away from including a RAID level in the vdisk attributes since RAID redundancy or striping is a property of the volume upon which the file system is built, rather than an attribute of the vdisk file. App. Br. 34-35. Appellants add that the Examiner's reliance on Busser and Hashemi is flawed since (1) Busser's RAID level information is not stored in a file, but rather accessed directly, and (2) Hashemi does not suggest storing striping information with virtual direct access storage device data. App. Br. 35-36; Reply Br. 10-11. The issues before us, then, are as follows:

## ISSUES

(1) Under § 103, has the Examiner erred by finding that Chen '509, Busser, and Hashemi collectively would have taught or suggested the specified internal organization includes (1) a RAID level, and (2) a striping pattern as recited in claims 28 and 29, respectively?

(2) Is the Examiner's reason to combine the teachings of these references supported by articulated reasoning with some rational underpinning to justify the Examiner's obviousness conclusion?

## ADDITIONAL FINDINGS OF FACT (FF)

19. Busser's system uses on-disk metadata to assist in recovering a RAID system 10. Each RAID system disk 14, 16, 18, 20, 22 contains stored data 90 which includes (1) disk metadata 100 which the controller 26 uses to

assist RAID operation and management, and (2) customer data 200. Busser, ¶¶ 0029-30; Figs. 1-2.

20. Busser's disk metadata 100 includes data fields that are read and periodically updated by the controller 26. These fields include RAID level 132 that corresponds to the particular RAID architecture used by a particular array. Busser, ¶¶ 0030, 0032; Fig. 2.

21. Hashemi discloses a RAID data storage system 200 that stripes data across an array of storage devices. An array controller 290 maintains the mapping between each virtual device and disk drives in the array. Hashemi, col. 5, l. 33 – col. 6, l. 30; Figs. 2-3E.

22. Hashemi notes that each LUN in array A300 can be further partitioned or concatenated with other LUNs to form smaller or larger LUNs for defining other RAID attributes like those of striping and/or redundancy. Hashemi, col. 10, l. 66 – col. 11, l. 5; Fig. 3A.

## ANALYSIS

We will sustain the Examiner's rejection of claims 28 and 29 which call for the specified internal organization to include (1) a RAID level, and (2) a striping pattern, respectively.

At the outset, we note that while the Examiner's rejected claims 28 and 29 collectively over the teachings of Chen '509, Busser, and Hashemi, the Examiner cited Busser to cure the perceived deficiencies of Chen '509 regarding claim 28. Ans. 12. The Examiner likewise cited Hashemi to cure the perceived deficiencies of Chen '509 regarding claim 29. Ans. 13. Since claims 28 and 29 each depend from claim 27, the teachings of Busser and

Hashemi are respectively cumulative regarding the limitations of claims 28 and 29 for which they were not relied upon—an error that we nevertheless deem harmless.

Turning to the rejection, we find no error in the Examiner's reliance on Busser for at least suggesting a RAID level internal organization as recited in claim 28. Each RAID system disk in Busser contains stored data that includes disk metadata which is used to assist in RAID operation and management—metadata that includes a RAID level field. FF 19-20. Despite Appellants' arguments to the contrary (App. Br. 34-36), we see no reason why such information could not be included as part of the specified internal organization in connection with the virtualization functionality of Chen '509 and Rajan.

Even assuming, without deciding, that the RAID redundancy or striping is a property of the file system volume in Chen '509 and Rajan as Appellants contend (App. Br. 34), we disagree that this property teaches away from providing RAID level information in connection with the specified internal organization in connection with the disclosed virtualization functionality noted previously. *See* FF 6-13. Nor have Appellants provided any persuasive evidence on this record proving that providing such RAID level information in this manner would have been beyond the level of ordinarily skilled artisans. Rather, we find this enhancement is tantamount to the predictable use of prior art elements according to their established functions—an obvious improvement. *See KSR Int'l Co. v. Teleflex, Inc.*, 550 U.S. 398, 417 (2007).

We reach a similar conclusion regarding the striping pattern of claim 29. In short, we see no error in the Examiner's reliance on Hashemi for this feature (Ans. 13), particularly since Hashemi explicitly teaches striping data across an array of devices, and that such data could be used in connection with LUNs for a particular array. FF 21-22. Despite Appellants' arguments to the contrary (App. Br. 36; Reply Br. 10-11), we find including striping pattern information in connection with the previously-discussed specified internal organization is tantamount to the predictable use of prior art elements according to their established functions—an obvious improvement. *See KSR*, 550 U.S. at 417. We therefore find the Examiner's reason to combine the teachings of the cited references supported by articulated reasoning with some rational underpinning to justify the Examiner's obviousness conclusion.

We are therefore not persuaded that the Examiner erred in rejecting claims 28 and 29.

## THE OBVIOUSNESS REJECTION OVER CHEN '509 AND CHEN '553

We will also sustain the Examiner's rejection of claim 33 over Chen '509 and Chen '553 (Ans. 13-14). Although Appellants contend that Chen '553 does not provide a motivation to include the recited internal organization specification in the file that stores the virtual direct access storage device attributes and data (App. Br. 37), the Examiner did not cite Chen '553 for that reason. Rather, the Examiner cited Chen '553 merely to show that backing up files over a network was well known, and therefore providing such a recited IP replication facility in view of the collective teachings of the cited references would have been obvious (Ans. 13-14)—a

position that we find reasonable. Moreover, we are unconvinced of error in the Examiner's reliance on Chen '509 for the reasons noted previously. The rejection of claim 33 is therefore sustained. And despite Appellants' arguments (App. Br. 37-38), we likewise are unpersuaded of error in the Examiner's obviousness rejection of claims 40 and 45-48 over Chen '509 and Chen '553 (Ans. 14-16)—a position that we find reasonable.

We are therefore not persuaded that the Examiner erred in rejecting claims 33, 40, and 45-48.

### THE OBVIOUSNESS REJECTION OVER CHEN '509, CHEN '553, BUSSER, AND HASHEMI

We will also sustain the Examiner's rejection of claims 41 and 42 over Chen '509, Chen '553, Busser, and Hashemi (Ans. 17-18). Despite Appellants' arguments to the contrary (App. Br. 38-39), we are unpersuaded of error in the Examiner's rejection essentially for the reasons indicated previously. The rejection is therefore sustained.

### THE OBVIOUSNESS REJECTION OVER CHEN '509, CHEN '553, AND BOLOSKY

Regarding claims 31, 32, 43, 44, and 49, the Examiner finds that the cited prior art discloses every recited feature except for the client initiating replication over a second concurrent TCP/IP connection. The Examiner, however, cites Bolosky for this feature in concluding the claim would have been obvious. Ans. 18.

Appellants argue that Bolosky does not provide a motivation to include the recited internal organization specification in the file that stores the virtual direct access storage device data (App. Br. 40-41). Appellants

add that in light of the substantially different environment of Bolosky's media server, there is insufficient motivation to modify the Chen references in light of Bolosky to arrive at the claimed invention as the Examiner proposes. App. Br. 41-44. The issues before us, then, are as follows:

## ISSUES

1. Under § 103, has the Examiner erred in rejecting claims 31, 32, 43, and 44 by finding that Chen '509, Chen '553, and Bolosky collectively would have taught or suggested programming the network file server to permit (1) a client to write new data to the virtual direct access storage device over a first TCP/IP connection, and (2) initiate copying of the file upon receipt of a command from the client over a second TCP/IP connection?

2. Is the Examiner's reason to combine the teachings of these references supported by articulated reasoning with some rational underpinning to justify the Examiner's obviousness conclusion?

## ADDITIONAL FINDINGS OF FACT (FF)

23. Bolosky discloses a computer system 12 connected to (1) controller 14 via control link 18, and (2) media storage 16 via data funnel 20. The control link facilitates messages to be passed between the controller and computer system, and the data funnel allows data to be transferred between the viewer and the controller. Both the control link and data funnel use the TCP/IP protocol. Bolosky, ¶¶ 0020-21; Fig. 1.

## ANALYSIS

We will sustain the Examiner's rejection of claims 31, 32, 43, 44, and 49 over Chen '509, Chen '553, and Bolosky. First, Appellants' contention that Bolosky does not provide a motivation to include the recited internal organization specification in the file that stores the virtual direct access storage device data (App. Br. 40-41, 44) is inapposite, for the Examiner did not cite Bolosky for that reason. Rather, the Examiner cited Bolosky merely to show that a client initiating a command over a second concurrent TCP/IP connection was well known (Ans. 18)—a position that we find reasonable given Bolosky's use of two different communication links that each use TCP/IP protocols. FF 23.

In light of this teaching, we find no error in the Examiner's position that using different TCP/IP connections to respectively (1) write data to the virtual direct access storage device, and (2) communicate client commands to initiate file copying would have been obvious. Despite Appellants' arguments to the contrary (App. Br. 40-44), we find that this enhancement to the arrangement disclosed by the Chen references is tantamount to the predictable use of prior art elements according to their established functions—an obvious improvement. *See KSR*, 550 U.S. at 417.

We are therefore not persuaded that the Examiner erred in rejecting claims 31, 32, 43, 44, and 49.

## CONCLUSION

Under § 102, the Examiner did not err in rejecting claims 27, 30, and 34-36. Under § 103, the Examiner did not err in rejecting claims 28, 29, 31-33, and 40-49, but erred in rejecting claims 7-11 and 19-23.

ORDER

The Examiner's decision rejecting claims 7-11, 19-23, 27-36, and 40-49 is affirmed-in-part.

No time period for taking any subsequent action in connection with this appeal may be extended under 37 C.F.R. § 1.136(a)(1)(iv).

AFFIRMED-IN-PART

pgc

RICHARD AUCHTERLONIE
NOVAK DRUCE & QUIGG, LLP
1000 LOUISIANA
53RD FLOOR
HOUSTON, TX 77002